

Transverse Velocity Field Measurement of Solar High-resolution Images Based on Unsupervised Deep Learning Postprint

Authors: Zhen-Hong Shang, Long Chen, Zhen-Ping Qiang, Yi Bi and Run-Xin Li

Date: 2025-04-28T11:57:59+00:00

Abstract

Measuring the transverse velocity field in high-resolution solar images is essential for understanding solar dynamics. This paper introduces an innovative unsupervised deep learning optical flow model designed to calculate the transverse velocity field, addressing the challenges of missing optical flow labels and the limited accuracy of velocity field measurements in high-resolution solar images. The proposed method converts the transverse velocity field computation problem into an optical flow computation problem, using two forward propagations of features to get rid of the reliance on optical flow labels. Additionally, it reduces the impact of the “Brightness Consistency” constraint on optical flow accuracy by identifying and handling optical flow outliers. We apply this method to compute the transverse velocity fields of high-resolution solar image sequences from the H α and TiO bands, observed by the New Vacuum Solar Telescope. Comparative experiments with several well-established optical flow methods, including those based on supervised deep learning models, show that our approach outperforms the comparison methods according to key evaluation metrics such as Residual Map Mean, Residual Map Variance, Cross Correlation, and Structural Similarity Index Measure. Moreover, since optical flow captures the fundamental motion information in image sequences, the proposed method can be applied to a variety of research areas, including solar image registration, sequence alignment, image super-resolution, magnetic field calibration, and solar activity forecasting. The code is available at <https://github.com/jackie-williamm/Transverse-Velocity-Field-Measurement-of-Solar-High-Resolution-Images>.

Full Text

Preamble

Research in Astronomy and Astrophysics, 25:035020 (13pp), 2025 March

© 2025. National Astronomical Observatories, CAS and IOP Publishing Ltd. All rights reserved, including for text and data mining, AI training, and similar technologies. Printed in China. <https://doi.org/10.1088/1674-4527/adbc38> CSTR: 32081.14.RAA.adbc38

Transverse Velocity Field Measurement of Solar High-resolution Images Based on Unsupervised Deep Learning

Zhen-Hong Shang¹², Long Chen¹, Zhen-Ping Qiang³, Yi Bi⁴, and Run-Xin Li¹²

¹ Faculty of Information Engineering and Automation, Kunming University of Science and Technology, Kunming 650500, China; rxli@kust.edu.cn ² Yunnan Key Laboratory of Artificial Intelligence, Kunming University of Science and Technology, Kunming 650500, China ³ College of Big Data and Intelligent Engineering, Southwest Forestry University, Kunming 650224, China ⁴ Yunnan Observatories, Chinese Academy of Sciences, Kunming 650216, China

Received 2025 January 5; revised 2025 February 22; accepted 2025 February 28; published 2025 March 24

Abstract

Measuring the transverse velocity field in high-resolution solar images is essential for understanding solar dynamics. This paper introduces an innovative unsupervised deep learning optical flow model designed to calculate the transverse velocity field, addressing the challenges of missing optical flow labels and the limited accuracy of velocity field measurements in high-resolution solar images. The proposed method converts the transverse velocity field computation problem into an optical flow computation problem, using two forward propagations of features to eliminate reliance on optical flow labels. Additionally, it reduces the impact of the “Brightness Consistency” constraint on optical flow accuracy by identifying and handling optical flow outliers. We apply this method to compute the transverse velocity fields of high-resolution solar image sequences from the H α and TiO bands, observed by the New Vacuum Solar Telescope. Comparative experiments with several well-established optical flow methods, including those based on supervised deep learning models, show that our approach outperforms comparison methods according to key evaluation metrics such as Residual Map Mean, Residual Map Variance, Cross Correlation, and Structural Similarity Index Measure. Moreover, since optical flow captures fundamental motion information in image sequences, the proposed method can be applied to various research areas, including solar image registration, sequence alignment,

image super-resolution, magnetic field calibration, and solar activity forecasting. The code is available at <https://github.com/jackie-williamm/Transverse-Velocity-Field-Measurement-of-Solar-High-Resolution-Images>.

Key words: methods: data analysis – techniques: image processing – Sun: fundamental parameters

1. Introduction

The velocity of material motion is a critical factor in studying solar eruptive activity and plays a key role in uncovering the underlying mechanisms of solar activity. Verma et al. (2013) systematically applied local correlation tracking (LCT) to analyze transverse motions in continuous solar images, building upon prior foundational work by November & Simon (1988) and Fisher & Welsch (2008). They demonstrated that LCT could accurately reconstruct certain properties of solar granulation, including the relationship between mean flow speed and averaging time, by comparing LCT velocities with those derived from simulated solar granulation data. Dumitru (2015) provided insights into the evolution of active region (AR) 10930, which produced numerous flares. Using LCT, they computed plasma velocities, offering a new perspective on the complex dynamics of this AR. Morgan & Hutton (2018) developed a method for calculating vector velocity fields based on propagating disturbances, revealing the presence of such disturbances in the corona. Additionally, Tremblay et al. (2021) utilized the advanced DeepVelU model to infer flow fields in the upper convection zone, chromospheric basement, and photosphere, demonstrating that deep learning is an effective tool for predicting solar activity. As telescope technology has advanced, researchers have collected vast amounts of high-resolution solar observation data, increasing the demand for accurate transverse velocity field measurements.

Approaches for measuring the solar transverse velocity field can be broadly categorized into conventional methods and deep learning-based methods, with the latter further divided into supervised and unsupervised techniques. LCT (November & Simon 1988) is one of the most widely used conventional methods in solar physics. It estimates the relative displacement of small sub-regions centered on specific pixels by maximizing the cross-correlation between consecutive images. To address some of LCT's limitations, the Fourier transform-based local correlation tracking (FLCT) method, proposed by Fisher & Welsch (2008), employs a Fourier-based approach to locate the peak of local correlation with subpixel precision. This improves both the accuracy and efficiency of tracking fluid motion.

Despite these advancements, FLCT shares similar challenges with LCT when applied to high-cadence data or fine-scale structures, particularly in dynamic regions with rapid intensity changes. Advances in deep learning have spurred interest in developing new approaches for measuring the solar transverse velocity field. DeepVel (Ramos et al. 2017), a fully convolutional neural network,

was designed to overcome the limitations of conventional methods. Trained on simulated velocity fields of the quiet Sun, DeepVel demonstrates potential in detecting small-scale vortices and fine structures. However, differences between simulated velocity fields and real solar velocity fields, particularly in high-resolution solar image sequences, may introduce challenges in accurately calculating velocity fields from such observations.

Directly measuring velocity vectors in solar observations is challenging due to the three-dimensional nature of solar activity and the limitations of current observational techniques. To address this, researchers commonly derive motion vectors from sequential images and compute the velocity field using temporal information. This transforms the problem of determining the transverse velocity field into one of estimating the optical flow field. Optical flow refers to the instantaneous velocity of pixel motion in the image plane of a moving object (Barron et al. 1992). By assuming the constancy of pixel brightness between consecutive frames, the Horn–Schunck (HS) algorithm (Horn & Schunck 1981) resolves the aperture problem in optical flow calculation and generates a dense optical flow field. Subsequently, the Lucas–Kanade (LK) algorithm (Lucas & Kanade 1981) assumes constant magnitude and direction of flow within a local region. Using the least squares method, it computes sparse optical flow fields by solving pixel flow vectors. Both methods rely on the brightness constancy assumption to formulate the optical flow energy equation, categorizing them as differential optical flow methods.

In the field of deep learning-based optical flow estimation for image sequences, PWCNet (Sun et al. 2018) has achieved notable success on several simulation datasets by integrating traditional techniques with supervised deep learning. The model excels in handling large displacements by prioritizing coarse-scale optical flow but struggles with small targets and minor displacements. To overcome this limitation, Shang et al. (2023) proposed PWCNet-S, specifically designed to measure the transverse velocity field in high-resolution solar images. PWCNet-S effectively improves the accuracy of optical flow estimation for small targets. However, due to the lack of true optical flow labels, the authors relied on synthetic labels generated by Demons (Liu et al. 2018), which introduced inaccuracies that affected the model’s overall performance. The quality of these labels significantly influences the accuracy of supervised deep learning models. In real-world applications, generating high-quality optical flow labels remains a challenge.

To address the limitations posed by missing labels and the discrepancy between synthetic and real data, unsupervised deep learning algorithms for optical flow have gained traction in recent years. These methods construct objective functions based on the assumptions of brightness constancy and local smoothness (Yu et al. 2016). For instance, UFlow (Jonschkowski et al. 2020) incorporates augmentation regularization during each iteration step to enhance model performance iteratively, while UPFlow (Luo et al. 2021) improves PWCNet’s up-sampling unit and introduces a more effective pyramid distillation loss, further

advancing the state of the art in unsupervised optical flow estimation. Despite these advancements, unsupervised methods still lag behind supervised approaches in terms of accuracy, particularly in scenarios with complex motion patterns or challenging illumination conditions.

High-resolution solar image sequences exhibit abundant self-similar structures with significant intensity variations, making them distinct from typical natural image sequences. The Sun's complex background further complicates the calculation of optical flow. Several challenges arise when processing high-resolution solar images with current deep-learning-based optical flow models: (1) Supervised deep learning optical flow models primarily rely on synthetic or simulated data from natural images, which differ significantly from real solar images (Torralba & Efros 2011; Mayer et al. 2018). Additionally, there is a lack of synthetic datasets that accurately capture the characteristics of solar images. (2) Unsupervised deep learning optical flow models often struggle to accurately compute optical flow in high-resolution solar images. For example, when a structure disappears, the model may incorrectly map multiple pixel points from the structure to a single location. Similarly, when a structure expands, pixel points resembling the structure may appear, leading to incorrect pixel matching. Furthermore, the brightness constancy assumption (Horn & Schunck 1981; Lucas & Kanade 1981) inherent in existing unsupervised models is frequently violated due to significant intensity changes observed in high-resolution images of solar active regions, reducing the precision of optical flow estimates. (3) While current optical flow methods perform well with large target displacements, there remains room for improvement in accurately capturing small structures and complex backgrounds in high-resolution solar images.

To address these challenges, this paper introduces an innovative unsupervised deep-learning optical flow model, which redefines the transverse velocity field computation as an optical flow estimation problem. By leveraging two forward propagations of features, the proposed method eliminates reliance on optical flow labels, while effectively mitigating the limitations imposed by the brightness consistency constraint through the identification and handling of optical flow outliers. This approach not only enhances the accuracy of optical flow computation but also provides a novel framework for precise transverse velocity field measurements in high-resolution solar images.

2. Data

The data used in this study were obtained from the New Vacuum Solar Telescope (NVST) at the Fuxian Lake Solar Observatory, Yunnan Observatories, Chinese Academy of Sciences (Liu et al. 2014). Two typical NVST observation bands were selected, resulting in two datasets. The datasets focus on active regions because their rich dynamic phenomena and complex structures, as observed in high-resolution image sequences, present significant challenges for velocity field estimation. In contrast, quiet-Sun regions and coronal holes, characterized by slower and more uniform flows, were excluded to prioritize scenarios where

precise velocity field estimation is both critical and particularly challenging in the context of high-resolution solar surface observations.

Dataset 1 consists of chromospheric line-center images observed by NVST in the $H\alpha$ (6563 Å) band. These images, acquired in AR 12178 between 07:26:20 and 09:28:51 UT on 2014 October 3, predominantly depict a filament eruption. The filament, which consists of two segments (left and right), is smoothly connected at a bending point (highlighted by the red dotted box in Figure 1(a)). The eruption began at 07:26 UT. As described by Xue et al. (2016), two reconnection events were observed during the eruption. The first occurred in the right section of the filament, where the eruption caused the two legs of the filament to move closer, eventually leading to reconnection. The second event occurred at the tail of the right filament (highlighted in the blue dashed circle in the lower left corner of Figure 1(a)), where it reconnected with chromospheric fibrils during the eruption. This dataset comprises 607 images with a 12 s cadence. The field of view is 540×590 pixels, with a pixel scale of 0.13 pixel^{-1} . In the chromosphere, a typical transverse velocity of 1 km s^{-1} corresponds to a displacement of 0.13 pixels. The image sequence shows subpixel and superpixel displacements across the entire field. Figure 1(a) displays a high-resolution $H\alpha$ image from dataset 1, where feature structures such as filaments and fibrils exhibit significant self-similarity.

Dataset 2 consists of images of the photosphere observed by NVST in the TiO band on 2022 May 29, from 01:28:04 to 04:32:07 UT. This dataset contains 301 images with a 36 s cadence. The field of view is 384×448 pixels, with a pixel scale of 0.052 pixel^{-1} . In the photosphere, a typical transverse velocity of 1 km s^{-1} corresponds to a displacement of 0.95 pixels. The velocity field exhibits subpixel and superpixel displacements across the complete field. Figure 1(b) showcases characteristic structures covering the umbra and penumbra of sunspots and granules.

Unlike supervised learning-based optical flow computation methods, the proposed approach eliminates the need for generating optical flow labels, which are challenging to obtain for real high-resolution solar images. To maintain distinct image distributions between the training and testing datasets, dataset 1 is split into training and testing subsets at an approximate 5:1 ratio in chronological order. Each consecutive image pair in the sequence constitutes an image pair. Specifically, the first 500 image pairs from dataset 1 are designated for training, while the remaining 106 pairs are allocated for testing, resulting in a total of 606 image pairs. Importantly, all 300 image pairs from dataset 2 are exclusively assigned to the test set and are not utilized in network training. Given the distinct structural characteristics of dataset 1 and dataset 2, the inclusion of dataset 2 provides an effective assessment of the model's generalization ability. Compared to dataset 2 (photosphere), dataset 1 (chromosphere) exhibits greater complexity and encompasses a broader range of structures and activities, thereby enriching the model's training process.

3. Method

This paper proposes an innovative unsupervised deep learning optical flow model that does not require optical flow labels, specifically designed to tackle the challenges of calculating optical flow in high-resolution solar images. The model integrates temporal information from image sequences to compute the transverse velocity field of these high-resolution images. The overall structure of the model is visualized in Figure 2. For each image pair, two forward propagations are performed: the first computes the optical flow using the original images, while the second computes the optical flow after applying data augmentation (DA). The optical flow from the first propagation is referred to as the “first optical flow,” while the optical flow from the second propagation is termed the “second optical flow.”

The model’s backbone consists of feature encoders and optical flow decoders. The feature encoder uses convolutional layers with residual blocks (He et al. 2016) to extract features from consecutive frames I_n and I_{n+1} . The optical flow decoder processes features using techniques such as bilinear upsampling (Dosovitskiy et al. 2015), optical flow warping (Brox et al. 2004), and cost-volume (Hosni et al. 2012), generating an optical flow pyramid $w_i(x)$ at resolutions of $1/8$, $1/4$, $1/2$, and full scale. Here, x denotes the pixel position, and i indicates the pyramid layer. The loss function for the first forward propagation is computed using gradient loss, smoothing loss, and photometric consistency loss. By employing census-transformed (Meister et al. 2017) image pairings, the brightness consistency loss becomes more robust to significant brightness variations.

In implementing the unsupervised learning approach, unlike supervised optical flow computation methods, endpoint errors are excluded from the loss function during the first forward propagation. The original image and the optical flow from the first forward propagation are augmented to generate soft labels for the optical flow of the augmented image. The augmented image is then passed through the model in the second forward propagation (SFP), where the soft label supervises the optical flow produced by the SFP. This strategy preserves the structural integrity of the solar images during DA while enhancing the model’s generalization capabilities. Additionally, bidirectional optical flow is used to identify pixels that deviate from the brightness consistency assumptions, and these pixels are excluded from the loss calculation to enhance model accuracy.

3.1. Unsupervised Learning Strategies

In supervised models, the supervisory effect of optical flow labels is typically achieved by calculating endpoint errors in the loss function. To enable unsupervised learning, we exclude the endpoint error calculation from the loss function during the first forward propagation. For each pair of consecutive images used in training, the model undergoes two forward propagations. As illustrated in Figure 2, the upper channel corresponds to the first forward propagation, while

the lower channel corresponds to the second. The original image pair is used in the first forward propagation, while the data-augmented image pair is used in the second. After the first forward propagation, the model computes the first optical flow $w_{1st}(x)$ between the original image pairs I_n and I_{n+1} . DA is then applied to I_n , I_{n+1} , and $w_{1st}(x)$, involving operations such as stretching, cropping, and erasing, resulting in the augmented image pairs \bar{I}_n and \bar{I}_{n+1} , and the corresponding soft labels \bar{w} for optical flow. In the SFP, \bar{I}_n and \bar{I}_{n+1} are input into the model to compute the second optical flow $w_{2nd}(x)$. The endpoint error between $w_{2nd}(x)$ and the soft labels \bar{w} is then computed in the loss function. The model is backpropagated by summing the loss functions from both forward propagations.

In our experiments, we observed that directly removing endpoint errors from the loss function during the first forward propagation led to poor model performance. Without endpoint errors, the model lacks guidance from optical flow labels, and the application of DA introduces complex scenarios in which the unsupervised loss becomes unreliable, significantly affecting the results (Liu et al. 2020). Excessive DA can distort the structure and characteristics of high-resolution solar images, particularly those containing small-scale features, making it difficult for the model to learn accurate optical flow. To address this issue, we employ two forward propagations to achieve self-supervision, which improves both the reliability and generalization of the unsupervised model by transforming the computation based on the original data.

3.2. Outlier Detection

When calculating optical flow between two consecutive images, I_n and I_{n+1} , pixels that are present in I_n but absent in I_{n+1} are referred to as “vanishing” pixels. Once they vanish, there is no corresponding pixel in the subsequent frame. These pixels do not satisfy the brightness consistency assumption in unsupervised optical flow models and can only be constrained by the smoothing assumption. This limitation may cause multiple pixels from I_n to be incorrectly mapped to the same location in I_{n+1} . Conversely, pixels that appear in I_{n+1} but are absent in I_n are termed “appearing” pixels. Since these newly appearing pixels have no corresponding points in the previous frame, unsupervised deep learning methods that rely solely on brightness and smoothing assumptions are prone to incorrect matches between locations in I_n and I_{n+1} , particularly when similar structures emerge. This mismatch can lead to erroneous optical flow results.

In dataset 1, as the eruption began, the filament initially moved downward and then shifted toward the lower right. The eastern part of the filament interacted with the western segment, as indicated by the blue dotted line in Figure 3. Subsequently, additional threads appeared around the original filament, also highlighted by the blue dotted line. In the optical flow model, a similar structure emerged around these filaments, corresponding to the “appearing” pixels. This dynamic behavior violates the brightness consistency assumption, upon

which traditional optical flow algorithms rely. These algorithms assume that brightness variations in the same physical region should be relatively smooth. Consequently, brightness-based optical flow computation methods fail or yield significant errors in solar image sequences with such intense activities.

To address the aforementioned issue, we propose a detection method based on the forward-backward consistency assumption (Meister et al. 2017). This assumption asserts that, for corresponding pixel points, the forward and backward optical flow vectors should have equal magnitudes but opposite directions. However, the model may make errors when calculating the optical flow of certain pixels due to the presence of “disappearing” and “appearing” pixels. To mitigate this, the model computes the forward optical flow $w_f(x)$ from image I_n to I_{n+1} and the backward optical flow $w_b(x)$ from I_{n+1} to I_n during forward propagation. By considering both the direction and magnitude of the optical flow field, a position x is classified as a non-special pixel if the total change in magnitude at that position is smaller than the sum of the weighted magnitudes of the forward and backward flows. In this case, the corresponding position in the mapping matrix o_f^x is set to 1; otherwise, it is set to 0.

In this study, we integrate the locations of non-special pixels into the loss function using the mapping matrix o_f^x . This ensures that the model prioritizes these pixels during the optical flow computation, thereby enhancing overall accuracy. The mapping matrix for non-special pixels is defined in Equation (1):

$$o_f^x = \begin{cases} 1 & \text{if } \|w_f(x) + w_b(x + w_f(x))\| < \alpha_1 \|w_f(x)\| + \alpha_2 \|w_b(x + w_f(x))\| \\ 0 & \text{otherwise} \end{cases}$$

The left side of the equation represents the agreement between the results of two optical flows, where $w_b(x + w_f(x))$ denotes the reversed forward optical flow. The right side of the equation defines a threshold function, where α_1 and α_2 are threshold parameters that control the allowable mismatch range for optical flow. We referenced the UnFlow (Meister et al. 2017) and set $\alpha_1 = 0.01$, $\alpha_2 = 0.5$. The mapping matrix o_f^x has dimensions of $H \times W$, where H and W denote the height and width of the image, respectively, and indicate the locations of special pixels. The mapping matrix o_b^x is computed in the same manner.

3.3. Loss Function

As illustrated in Figure 2, we adopt a novel training strategy that involves two forward propagations of the model for each image pair during training. In both forward propagations, the model computes the photometric loss \mathcal{L}_{photo} , gradient loss $\mathcal{L}_{gradient}$, and first-order smoothing loss \mathcal{L}_{smooth} to ensure that the model is able to take into account pixel photometry, gradient variation, and smoothness in computing the optical flow as described in Equations (6), (8) and (9). The soft labeling of the optical flow, generated during the first forward propagation,

provides a crucial guide for the second propagation, despite its limited accuracy. In the SFP, the endpoint error loss \mathcal{L}_{epe} is added to the loss function. Thus, after two forward passes, the model's total loss function combines the first forward propagation loss and the weighted SFP loss, as presented in Equation (2):

$$\mathcal{L}_{total} = \mathcal{L}_{1st} + \omega_{2nd}\mathcal{L}_{2nd}$$

Here, \mathcal{L}_{1st} denotes the loss from the first forward propagation, while \mathcal{L}_{2nd} represents the loss from the SFP. Since the second forward pass uses data-augmented images to compute optical flow, it is prone to generating a large number of erroneous optical flow estimates, which can lead to an excessively large \mathcal{L}_{2nd} . To mitigate this issue and ensure a proper adjustment effect from the SFP, we set ω_{2nd} to 0.01.

3.3.1. Loss Function for Two Forward Propagations The model's two forward propagations each generate four hierarchical levels of optical flow, corresponding to different scales, ranging from coarse to fine. The loss for each level is calculated and denoted as \mathcal{L}_i , where $i = 1, 2, 3, 4$. We assign smaller weight factors to the loss functions computed at the coarse scales and larger weight factors to those computed at the finer scales. This weighting scheme enables the network to converge more rapidly at the coarse scale, providing a strong initialization for subsequent refinement at finer scales, as shown in Equation (3):

$$\mathcal{L}_{1st} = \sum_{i=1}^4 \gamma^{i-1} \mathcal{L}_i$$

where γ^{i-1} represents the weighting factor, with γ set to 0.5 in the implementation. This approach allows our model to effectively capture motion information across different scales and produce high-quality optical flow results.

The loss function for each layer of forward propagation is given in Equation (4), with the following parameter settings: $\omega_{photo} = 0.5$, $\omega_{gradient} = 0.5$, $\omega_{smooth} = 0.005$.

$$\mathcal{L}_i = \omega_{photo}\mathcal{L}_{photo} + \omega_{gradient}\mathcal{L}_{gradient} + \omega_{smooth}\mathcal{L}_{smooth}$$

The loss function for each layer of the SFP is given in Equation (5):

$$\mathcal{L}_i^{SFP} = \omega_{photo}\mathcal{L}_{photo} + \omega_{gradient}\mathcal{L}_{gradient} + \omega_{smooth}\mathcal{L}_{smooth} + \omega_{epe}\mathcal{L}_{epe}$$

where \mathcal{L}_{epe} is computed using optical flow soft labeling and secondary optical flow calculation. The weight for the endpoint error loss is set to $\omega_{epe} = 1$, while the remaining loss weights are the same as those in the first forward propagation.

3.3.2. Photometric, Gradient, Smoothing, and Endpoint Error Loss Functions

Photometric consistency loss generally assumes that the brightness of a target pixel remains nearly constant across consecutive image frames, which holds true for many optical flow tasks in natural scenes. However, when processing high-resolution solar images, significant brightness variations can occur, leading to the failure of the photometric consistency loss function (Vogel et al. 2013). As shown in Figure 4, the filament eruption induces substantial brightness changes in the region highlighted by the red box. To overcome this challenge, our photometric consistency loss function computes luminosity consistency using image pairs after applying the ternary census transform (Stein 2004). This method compensates for both additive and multiplicative illumination changes, as well as gamma variations, thereby providing a more robust assumption of photometric constancy in real images (Hafner et al. 2013).

The photometric consistency loss function is defined in Equation (6):

$$\mathcal{L}_{photo} = \frac{1}{|o_f|} \sum_{x \in o_f} \Psi(I_n(x) - I_{n+1}(x + w_f(x))) + \frac{1}{|o_b|} \sum_{x \in o_b} \Psi(I_{n+1}(x) - I_n(x + w_b(x))) + \lambda_p \sum_{x \notin o_f} \Psi(I_n(x) - I_{n+1}(x + w_f(x)))$$

This loss function is composed of normalized photometric loss terms for forward and backward warping of non-special pixels, and a penalty term for special pixels weighted by λ_p . Here, o_f and o_b are binary mapping matrices and optical flow fields for forward and backward directions respectively (introduced in Section 3.2). H and W represent the height and width of the image, respectively. The function $\Psi(\cdot)$ denotes the robust generalized Charbonnier penalty function, with $\alpha = 0.45$ and $\epsilon = 0.001$ (Sun et al. 2014). To measure photometric differences, the census transform function $c(\cdot)$ is employed, which computes the difference of census transforms (Stein 2004) between corresponding pixels x in I_n and I_{n+1} : $I'_{n+1} = I_{n+1}(x + w_f(x))$, where $c(\cdot)$ represents the census transform (Stein 2004), which converts each local intensity pattern into a bit string, ensuring robustness to illumination variations. The backward photometric loss is computed in the same way. The term λ_p ($\lambda_p = 10$) introduces a small constant photometric penalty for regions identified as special pixels, such as those corresponding to disappearing or appearing structures. These regions violate the photometric consistency assumption and are masked by $1 - o_f$ or $1 - o_b$. The penalty term discourages the model from converging to a trivial solution where all pixels become special, and compels the network to prioritize a trade-off between photometric consistency and special pixels. This mechanism forces the network to balance photometric consistency loss against outlier handling, thereby encouraging the learning of realistic optical flow motion and reliable occlusion regions, ultimately enhancing the quality and effectiveness of optical flow estimation in dynamic solar scenes.

The gradient loss ensures that the estimated optical flow accurately captures edges and fine details in high-resolution solar images by penalizing significant

gradient differences between warped and target images. In solar images, gradients are closely related to important structures, such as filaments (in H α) and granules (in TiO), and the optical flow should track structural motion while preserving edge and texture continuity. The gradient loss emphasizes structural preservation, complementing photometric loss. While photometric loss ensures brightness consistency, it alone is insufficient for capturing complex structural motion. Gradient loss directly addresses this by focusing on edges and details. Inspired by PWCNet-S (Shang et al. 2023), our gradient loss is distinct: it includes both forward and backward components and is computed specifically for non-special pixels. The formulation is:

$$\mathcal{L}_{gradient} = \frac{1}{|o_f|} \sum_{x \in o_f} \Psi(\nabla I_n(x) - \nabla I_{n+1}(x + w_f(x))) + \frac{1}{|o_b|} \sum_{x \in o_b} \Psi(\nabla I_{n+1}(x) - \nabla I_n(x + w_b(x)))$$

Similar to the photometric loss, the gradient loss components are normalized by the denominators $|o_f|$ and $|o_b|$, respectively, to prevent any bias due to the number of non-special pixels in the image. Furthermore, since the penalty for special pixels is already incorporated in the photometric loss and is reflected in the total loss, there is no need for additional penalties in the gradient loss.

The smoothness loss is implemented to enforce spatial coherence within the computed optical flow field, particularly in local regions. While solar atmospheric plasma motion is inherently complex, within the interiors of solar structures, motion typically exhibits continuity and gradual variation. This is physically motivated and aligns with the expected behavior of solar phenomena. Due to the higher likelihood of errors in the upper layers of the optical flow pyramid, which can affect the accuracy of the output optical flow (Jonschkowski et al. 2020), we only calculate the smoothness loss for the output optical flow. The smoothing loss is shown in Equation (9), incorporating a smoothness constraint with edge-aware weighting:

$$\mathcal{L}_{smooth} = \sum_x \left| \frac{\partial w(x)}{\partial x} \right|_1 e^{-\beta \|\partial I_n(x)\|_1} + \left| \frac{\partial w(x)}{\partial y} \right|_1 e^{-\beta \|\partial I_n(x)\|_1}$$

Here, $\|\partial(w(x))\|_1$ denotes the L1 norm of the spatial gradient of the optical flow field $w(x)$ at pixel location x . Crucially, edge-aware weighting terms, $e^{-\beta \|\partial I_n(x)\|_1}$, are incorporated to selectively reduce the influence of the smoothness constraint at structural edges. This design allows the model to flexibly accommodate motion discontinuities precisely where they are physically reliable—at structural boundaries—while rigorously enforcing overall smoothness within the structural interiors. In contrast to the photometric and gradient losses, which operate on image pixels and gradients respectively, the smoothness loss directly acts upon the optical flow field itself. Functionally, it serves as a regularization term, effectively suppressing noise and promoting the generation of a more physically reliable and stable optical flow solution.

Exclusively employed during the SFP, the endpoint error loss serves a critical role in our training strategy. Its primary objective is to leverage the “soft labels,” denoted as \bar{w} , to guide and supervise the optical flow estimation in the second pass. These soft labels are generated from the first forward propagation and subsequently augmented with DAs. The endpoint error loss directly quantifies the disparity between the optical flow computed in the second forward pass, $w_{2nd}(x)$, and these soft labels:

$$\mathcal{L}_{epe} = \frac{1}{N} \sum_x \|w_{2nd}(x) - \bar{w}(x)\|_2$$

In unsupervised optical flow computation for high-resolution solar images, where true optical flow ground truth is absent, endpoint error loss provides an effective self-supervision mechanism. By using the first forward propagation’s results as learning targets to guide the SFP, the model can iteratively refine its optical flow estimation capabilities. This self-supervision is critical for handling the complexity and diversity of solar structures and enhancing the performance of unsupervised optical flow models. Especially in solar active regions with complex and variable structural motions, this self-supervised learning approach enables the model to better adapt to these challenging scenarios.

4. Experiments

4.1. Experimental Environment

We utilized computer systems equipped with Intel Xeon Gold 6226R CPUs and NVIDIA RTX A6000 GPUs running the Windows 10 operating system, and established an experimental environment based on Python and PyTorch 1.7.1 (Paszke et al. 2019) to ensure efficient operation and accuracy.

4.2. Evaluation Metrics

We use the trained model to compute the forward optical flow field from image I_n to I_{n+1} , and then obtain the reconstructed image \hat{I}_{n+1} by warping I_n toward I_{n+1} using the optical flow field. In order to comprehensively evaluate the difference between the reconstructed image \hat{I}_{n+1} and the target image I_{n+1} , we use four metrics: structural similarity (Wang et al. 2004) (SSIM), Cross-Correlation (CC), Residual Map Mean (RMM), and Residual Map Variance (RMV). These metrics collectively reflect the degree of similarity between the reconstructed and target images, with smaller differences indicating more accurate optical flow field computations.

SSIM is calculated as expressed in Equation (11). Here, $\mu_{\hat{I}_{n+1}}$ and $\mu_{I_{n+1}}$ denote the mean of image \hat{I}_{n+1} and image I_{n+1} respectively, $\sigma_{\hat{I}_{n+1}}^2$ and $\sigma_{I_{n+1}}^2$ denote the variance of image \hat{I}_{n+1} and I_{n+1} respectively, $\sigma_{\hat{I}_{n+1}I_{n+1}}$ denotes the covariance

between image \hat{I}_{n+1} and image I_{n+1} , and c_1 and c_2 are constants. The SSIM value ranges from 0 to 1, and being closer to 1 indicates that the reconstructed image \hat{I}_{n+1} is more similar to the target image I_{n+1} in terms of structure, signifying higher accuracy for the calculated optical flow.

The specific calculation formula for the CC is given in Equation (12). The CC values range from 0 to 1, with values closer to 1 indicating higher accuracy in the optical flow computation, reflecting a stronger correlation between the reconstructed and target images in terms of pixel value distribution.

The RMM and RMV are calculated based on the residual maps of the two images, with the specific formulas provided in Equations (13) and (14). A smaller RMM indicates that the reconstructed image \hat{I}_{n+1} is closer to the target image I_{n+1} in terms of grayscale intensity. Likewise, a smaller RMV suggests fewer anomalous matches in \hat{I}_{n+1} , reflecting greater stability in the algorithm's matching capability.

$$\text{RMM} = \frac{1}{HW} \sum_x |\hat{I}_{n+1}(x) - I_{n+1}(x)|$$

$$\text{RMV} = \frac{1}{HW} \sum_x (\hat{I}_{n+1}(x) - I_{n+1}(x))^2$$

4.3. Experimental Results and Analysis

During the training process, we chose the AdamW (Loshchilov & Hutter 2017) optimizer to update the model parameters and clip the gradient to ensure that its value is in the range of $[-1, 1]$, thus avoiding the gradient explosion problem. We set the maximum learning rate to 0.001 and the learning rate weight decay to 0.0001. The batch size is set to 2 and the OneCycleLR (Smith & Topin 2019) learning rate tuning scheme is used.

4.3.1. Quantitative Performance Comparison and Ablation Experiments

We conducted a comprehensive series of ablation and comparison tests to evaluate the performance of the model. To ensure fairness and consistency, all comparison methods were trained using the same protocol and datasets. Table 1 provides detailed results comparing our model with current algorithms on datasets 1 and 2, including the supervised deep learning method PWCNet-S (Shang et al. 2023), unsupervised deep learning methods UFlow (Jonschkowski et al. 2020) and UPFlow (Luo et al. 2021), as well as traditional methods commonly used in astronomy such as FLCT (Fisher & Welsch 2008) and Demons (Liu et al. 2018).

The proposed method performs better than alternative approaches on the primary evaluation metrics for datasets 1 and 2, as indicated in Table 1. We concentrate on SSIM, a crucial metric for analyzing image structure similarity

and determining the accuracy of optical flow computation. The original image pairs in dataset 1 have an average SSIM of 0.8418 and an RMV of 45.8184. The average SSIM between the reconstructed images \hat{I}_{n+1} and the target images I_{n+1} is 0.9875 and the RMV drops to 5.0818 when the proposed approach is used to calculate the optical flow and produce the reconstructed image. Datasets 1 and 2 feature different solar structures (see Figure 1). The proposed approach produced positive outcomes on dataset 2 during the training phase, despite the fact that the model cannot see it. The average SSIM between the target image and the reconstructed image on dataset 2 is 0.9796 following the application of the proposed method, representing an improvement of 11.77% over the original image. The proposed method lowers the RMV by 89.9% and 77.5%, respectively, and increases the SSIM on dataset 1 and dataset 2 by 18.5% and 14.99%, respectively, as compared to FLCT (Fisher & Welsch 2008), a traditional method used frequently in astronomy. The accuracy of supervised methods is often higher than that of unsupervised methods in deep learning optical flow. The proposed approach outperforms supervised method PWCNet-S (Shang et al. 2023) with respect to SSIM on datasets 1 and 2 by 4.32% and 0.25%, respectively. In conclusion, the proposed approach not only enhances optical flow computation accuracy but also exhibits strong generalization capabilities.

To evaluate the efficacy of each component of the proposed method, we conducted comprehensive ablation experiments on datasets 1 and 2, and we quantitatively analyzed the results presented in Table 2. The components tested are as follows: (a) adding an SFP; (b) using outlier detection (OD) to exclude the computation of losses at particular pixel points; and (c) using an improved photometric loss function and smoothing loss function (ILF). The BASELINE (shown in the upper channel of Figure 2) is obtained by subtracting the EPE from the PWCNet-S model's loss function; at this stage, the original image is used for training. As can be shown in Table 2's first row, the model outperforms PWCNet-S on dataset 1 because the unsupervised learning method prevents inaccurate optical flow in the optical flow labels produced by Demons (Shang et al. 2023). However, the model's performance on dataset 2 leaves something to be desired, with an average SSIM between the reconstructed images \hat{I}_{n+1} and the target images I_{n+1} of only 0.9405.

In order to improve the model's generalization performance, we use the SFP in the model, as shown in the second row of Table 2. Due to the accumulation of incorrect optical flow in the first forward propagation in the SFP, the model's SSIM on dataset 1 is somewhat worse. On dataset 2, however, the SSIM is marginally improved and the RMV of the reconstructed image relative to the target image is decreased by 9.87%. These results highlight the positive effect of SFP in improving the generalization ability of the model. By including OD in the loss function, which eliminates the impact of special pixels, the model's performance is significantly enhanced, as can be seen in the third row of Table 2. RMV on datasets 1 and 2 is reduced by 27.98% and 7.9%, respectively, as a result of this change. It is feasible to reduce the impact of special pixels and

raise the model's accuracy through the removal of pixels that do not adhere to the photometric assumptions. Finally, when using the improved photometric loss function as well as the smoothing loss function, the performance of the model on both datasets 1 and 2 is somewhat improved, as shown in the last row of Table 2. SSIM improves on both datasets, whereas the RMV is reduced by 9.03% on dataset 1 and 9.58% on dataset 2. It is apparent that the significance of a well-designed loss function is to raise the accuracy of unsupervised learning. In conclusion, the proposed method effectively enhances the model's performance on high-resolution solar images and, by integrating particular data enhancements, also enhances the model's capacity for generalization.

4.3.2. Qualitative Performance Comparison Experiments Two consecutive images from the test set of dataset 1 were utilized in a visualization experiment to further demonstrate the effectiveness of the proposed method. The two consecutive images used for testing are shown in Figure 5, corresponding to the stage when the filament eruption had concluded. This figure highlights the complex background present in the high-resolution solar image, which contains many fine structures.

Figures 6–8 present a comparative analysis of the optical flow computed by different methods. As can be seen in the first row of Figure 6, we create the reconstructed image by warping Figure 5(a) to Figure 5(b) using the optical flow calculated by various methods. The residual image is then obtained by subtracting the anticipated image from the target image (Figure 5(b)), as indicated in the second row of Figure 6. A larger view of the red-boxed area of the second row of images can be seen in the third row of Figure 6. The higher the accuracy of the optical flow prediction, the fewer the residuals, and the smaller the disparity between the target and reconstructed images. The residual images show that the reconstructed images calculated by the traditional astronomical method Demons, the supervised deep learning method PWCNet-S, and the proposed method are closer to the target images, and the optical flow is more accurate. Compared to FLCT, a widely used method for velocity field measurements in astronomy, the proposed approach significantly reduces the disparity between the reconstructed and target images, leading to a substantial improvement in optical flow accuracy. Furthermore, when compared to the supervised deep learning optical flow method PWCNet-S, the proposed method achieves an even higher level of accuracy in optical flow estimation.

Figure 7 depicts the vector arrow maps of the optical flow computed by different methods within the blue-boxed region (50×50 pixels) of Figure 5, which was selected to better visualize the flow field, with arrow colors representing the magnitude of the speed. The results in the figure show that PWCNet-S and Demons produce similar optical flow maps, as the supervised method PWCNet-S uses optical flow generated by Demons as soft labels during training. In terms of the arrow directions, the proposed method is comparable to these two methods (as shown in the green dotted box area in Figure 7). However, the optical flow maps

generated by PWCNet-S and Demons exhibit large uniform regions, indicating a lack of detailed variability in these areas. In contrast, the proposed method produces optical flow maps with significantly more intricate details, demonstrating its superior ability to capture fine-grained motion features and subtle variations within the image. This highlights the effectiveness of the proposed method in computing fine structures in high-resolution solar images.

Figure 8 presents the optical flow visualization maps for the same blue-boxed region of Figure 5, with the colors indicating the magnitude and direction of the optical flow.

5. Conclusions

Measuring the transverse velocity field in high-resolution solar images is crucial for understanding the dynamic characteristics of the solar atmosphere, as well as for predicting solar activities. We propose a method for measuring the transverse velocity field based on an unsupervised deep learning optical flow model. This method addresses the challenges of missing optical flow labels and improves velocity field accuracy by leveraging two forward passes of feature propagation to eliminate reliance on optical flow labels. Additionally, it includes outlier detection to minimize the impact of the “Brightness Consistency” constraint on flow accuracy. Extensive experiments have demonstrated that the proposed method is highly effective for high-resolution solar images, such as those captured by the NVST in the $H\alpha$ and TiO bands. The computed optical flow fields offer substantial potential for various research applications.

Since unsupervised models do not rely on ground truth data, we plan to use additional high-resolution solar images in future training to enhance the model’s generalization capabilities. Furthermore, because the motion of structures in solar images tends to be continuous, training the model using just two frames represents a conservative approach. In future work, we aim to incorporate sequences with three or more frames, which will provide more motion information and help further optimize the model’s performance.

Acknowledgments

We thank the NVST team for providing the data. This study is supported by the National Natural Science Foundation of China (NSFC, Grant Nos. 12063002 and 12163004).

ORCID iDs Zhen-Hong Shang <https://orcid.org/0000-0003-1280-5255>

References

Barron, J. L., Fleet, D. J., Beauchemin, S. S., & Burkitt, T. A. 1992, in Proc. 1992 IEEE Computer Society Conf. Computer Vision and Pattern Recognition (Los Alamitos, CA: IEEE Computer Society Press), 236

- Brox, T., Bruhn, A., Papenberger, N., & Weickert, J. 2004, in Computer Vision-ECCV 2004: 8th European Conf. Computer Vision, Prague (Berlin: Springer), 25
- Dosovitskiy, A., Fischer, P., Ilg, E., et al. 2015, in Proc. IEEE Int. Conf. Computer Vision (IEEE), 2758
- Dumitru, L. 2015, RoAJ, 25, 23
- Fisher, G. H., & Welsch, B. T. 2008, in ASP Conf. Ser. 383, Subsurface and Atmospheric Influences on Solar Activity, ed. R. Howe et al. (San Francisco, CA: ASP), 373
- Hafner, D., Demetz, O., Weickert, J., et al. 2013, in Scale Space and Variational Methods in Computer Vision: 4th Int. Conf., SSVM 2013 (Berlin: Springer), 210
- He, K., Zhang, X., Ren, S., & Sun, J. 2016, in Proc. IEEE Conf. Computer Vision and Pattern Recognition (Piscataway, NJ: IEEE), 770
- Horn, B. K., & Schunck, B. G. 1981, Artificial Intelligence, 17, 185
- Hosni, A., Rhemann, C., Bleyer, M., Rother, C., & Gelautz, M. 2012, ITPAM, 35, 504
- Jonschkowski, R., Stone, A., Barron, J. T., et al. 2020, in Computer Vision-ECCV 2020: 16th European Conf. (Berlin: Springer), 557
- Liu, H., Yang, Y., Shang, Z., & Li, R. 2018, AR&T, 15, 7
- Liu, L., Zhang, J., He, R., et al. 2020, in Proc. IEEE/CVF Conf. Computer Vision and Pattern Recognition (Piscataway, NJ: IEEE), 6489
- Liu, Z., Xu, J., Gu, B.-Z., et al. 2014, RAA, 14, 705
- Loshchilov, I., & Hutter, F. 2017, in Int. Conf. Learning Representations
- Lucas, B. D., & Kanade, T. 1981, in IJCAI'81: 7th Int. Joint Conf. Artificial Intelligence (San Francisco, CA: Morgan Kaufmann Publishers Inc.), 674
- Luo, K., Wang, C., Liu, S., et al. 2021, in Proc. IEEE/CVF Conf. Computer Vision and Pattern Recognition (Piscataway, NJ: IEEE), 1045
- Mayer, N., Ilg, E., Fischer, P., et al. 2018, International Journal of Computer Vision, 126, 942
- Meister, S., Hur, J., & Roth, S. 2017, in AAAI Conf. Artificial Intelligence, 32 (Association for the Advancement of Artificial Intelligence (AAAI))
- Morgan, H., & Hutton, J. 2018, ApJ, 853, 145
- November, L. J., & Simon, G. W. 1988, ApJ, 333, 427
- Paszke, A., Gross, S., Massa, F., et al. 2019, in Advances in Neural Information Processing Systems, 32 (Curran Associates, Inc.)

- Ramos, A. A., Requerey, I., & Vitas, N. 2017, A&A, 604, A11
- Shang, Z.-H., Mu, S.-Y., Ji, K.-F., & Qiang, Z.-P. 2023, RAA, 23, 065009
- Smith, L. N., & Topin, N. 2019, in Artificial Intelligence and Machine Learning for Multi-Domain Operations Applications, 11006 (SPIE), 36
- Stein, F. 2004, in Joint Pattern Recognition Symp. (Berlin: Springer), 79
- Sun, D., Roth, S., & Black, M. J. 2014, International Journal of Computer Vision, 106, 115
- Sun, D., Yang, X., Liu, M.-Y., & Kautz, J. 2018, in Proc. IEEE Conf. Computer Vision and Pattern Recognition, 8934 (Piscataway, NJ: IEEE), 8934
- Torralba, A., & Efros, A. A. 2011, in CVPR 2011 (Piscataway, NJ: IEEE), 1521
- Tremblay, B., Cossette, J.-F., Kazachenko, M. D., Charbonneau, P., & Vincent, A. 2021, JSWSC, 11, 9
- Verma, M., Steffen, M., & Denker, C. 2013, A&A, 555, A136
- Vogel, C., Roth, S., & Schindler, K. 2013, in Pattern Recognition: 35th German Conf., GCPR 2013, Proc. 35 (Berlin: Springer), 343
- Wang, Z., Bovik, A. C., Sheikh, H. R., & Simoncelli, E. P. 2004, ITIP, 13, 600
- Xue, Z., Yan, X., Cheng, X., et al. 2016, NatCo, 7, 11837
- Yu, J. J., Harley, A. W., & Derpanis, K. G. 2016, in Computer Vision-ECCV 2016 Workshops (Berlin: Springer), 3

Note: Figure translations are in progress. See original paper for figures.

Source: ChinaXiv — Machine translation. Verify with original.